

Oracle® CODASYL DBMS for OpenVMS

Table of Contents

<u>Oracle® CODASYL DBMS for OpenVMS</u>	1
<u>Release Notes</u>	2
<u>Preface</u>	5
<u>Purpose of This Manual</u>	6
<u>Intended Audience</u>	7
<u>Document Structure</u>	8
<u>Conventions</u>	9
<u>Chapter 1 Installing Oracle CODASYL DBMS</u>	10
<u>1.1 Oracle CODASYL DBMS on OpenVMS I64</u>	11
<u>1.2 Using Databases from Releases Earlier Than V7.0</u>	12
<u>1.3 Requirements</u>	13
<u>1.4 Installation of Oracle CODASYL DBMS Software</u>	14
<u>1.5 Documentation in Adobe Acrobat Format</u>	15
<u>Chapter 2 Enhancements Provided</u>	16
<u>2.1 Hardware Note on Alpha EV56 Requirement</u>	17
<u>2.2 Limited IEEE Floating Support in Data Definition Language</u>	18
<u>2.3 /NOSNAPSHOTS, /[NO]DATA FILE added to DBO/MOVE AREA</u>	19
<u>2.4 DBO/BACKUP/MULTI/GROUP SIZE Default Value Increased</u>	22
<u>2.5 The /LOCK TIMEOUT Qualifier Now Allows the Database Default</u>	23
<u>2.6 Global Statistics Sections for Better Performance</u>	24
<u>2.7 Maximum Buffer Size Increase</u>	25
<u>2.8 DBO/SHOW STATISTICS Global Sections Move to P2 Address Space</u>	26
<u>2.9 EXQUOTA Caused Inaccessible ALL</u>	27

Table of Contents

<u>2.10 Changes to Default Behavior in Oracle DBMS</u>	28
<u>2.11 Add option to pass values to /CONFIRM during RESTORE operation</u>	29
<u>2.12 Minimum number of allowed users increased</u>	31
Chapter 3Problems Corrected	32
<u>3.1 DBO/BACKUP[/MULTITHREAD/ONLINE Bugcheck at DIOBCK\$GET PAGE</u>	33
<u>3.2 DBO/SHOW STATISTICS Configuration File Issues</u>	34
<u>3.3 DBO/SHOW STATISTICS Access Violation on Hot Row Information Screen</u>	35
<u>3.4 Problems If a Full DBO/BACKUP Was Not Done After DBO/MOVE AREA</u>	36
<u>3.5 DBO Show Statistics Sometimes Bugchecks When Using Process Monitoring</u>	37
<u>3.6 Incorrect DBO/BACKUP/AFTER Truncate AIJ File Error Handling</u>	38
<u>3.7 DBO Show Statistics Sometimes Bugchecks on Row Cache Information Screen</u>	39
<u>3.8 Problem Writing Large TSN Values to Data and Snap Pages</u>	40
<u>3.9 DBO/RECOVER/ORDER AIJ FILES ERROR HANDLING PROBLEM WHEN ORDERING FILES</u>	41
<u>3.10 DBO/SHOW STATISTICS/CLUSTER Did Not Show Transaction Type</u>	43
<u>3.11 DBO Show Statistics not using the full screen size for details</u>	44
<u>3.12 Node Failure May Cause DBRs to Hang</u>	45
<u>3.13 DBO/SHOW STATISTICS (Transaction Duration (Total)) playback generates a bugcheck</u>	46
<u>3.14 %SMG-F-INVROW IN DBO/SHOW STAT USING OPTION WRITE REPORT GRAPH OR BOTH</u>	47
<u>3.15 Unable to Start DBO/SHOW STATISTICS During Cluster State Transition</u>	48
<u>3.16 DBO/BACKUP/ONLINE/NOQUIET did not back up the root log file entries</u>	49
<u>3.17 Unexpected bugcheck dump from DBO/SHOW STATISTICS/HOT STANDBY LOG</u>	50
Chapter 4Known Problems and Restrictions	51

Table of Contents

<u>4.1 Ensure No Processes Have DBMSHR Image Activated</u>	52
<u>4.2 Syntax for Prestart Transaction Timeout</u>	53
<u>4.3 Patch Required When Using VMS V8.3 or V8.2-1 on Itanium</u>	54
<u>4.4 /LOCK TIMEOUT Documentation Clarification</u>	55
<u>4.5 Unexpected RCS termination</u>	56
<u>4.6 Slight Relaxation Of VM\$MEM RESIDENT USER Requirement</u>	57
<u>4.7 Patch Required When Using VMS V8.3 and Dedicated CPU Lock Manager</u>	58
<u>4.8 VM\$MEM RESIDENT USER Rights Identifier Required</u>	59
<u>4.9 Features Not Yet Available for OpenVMS I64</u>	60
<u>4.10 Expect Additional Memory Consumption</u>	61
<u>4.11 ILINK-E-INVORINI Error on I64</u>	62
<u>4.12 Default Floating Point Format</u>	63
<u>4.13 SYSTEM-F-IN\$FMEM Fatal Error With SHARED SYSTEM MEMORY or LARGE MEMORY Enabled in Galaxy Environment</u>	64
<u>4.14 Oracle CODASYL DBMS and OpenVMS ODS-5 Volumes</u>	65
<u>4.15 Carryover Locks and NOWAIT Transaction Clarification</u>	66
<u>4.16 Both Application and Oracle CODASYL DBMS Using SY\$HIBER</u>	67
<u>4.17 Row Cache Not Allowed While Hot Standby Replication is Active</u>	69
<u>4.18 Exclusive Access Transactions May Deadlock with RCS Process</u>	70
<u>Chapter 5 New Features and Corrections in Previous Releases</u>	71

Oracle® CODASYL DBMS for OpenVMS

Release Notes

Release 7.3.1.0 for OpenVMS Alpha and
HP OpenVMS Industry Standard 64 for Integrity Servers

NOVEMBER 2013

Oracle CODASYL DBMS Release Notes, Release 7.3.1.0 for OpenVMS Alpha and OpenVMS I64

Copyright © 2013 Oracle Corporation. All rights reserved.

The Programs (which include both the software and documentation) contain proprietary information of Oracle Corporation; they are provided under a license agreement containing restrictions on use and disclosure and are also protected by copyright, patent and other intellectual and industrial property laws. Reverse engineering, disassembly or decompilation of the Programs, except to the extent required to obtain interoperability with other independently created software or as specified by law, is prohibited.

The information contained in this document is subject to change without notice. If you find any problems in the documentation, please report them to us in writing. Oracle Corporation does not warrant that this document is error-free. Except as may be expressly permitted in your license agreement for these Programs, no part of these Programs may be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose.

If the Programs are delivered to the U.S. Government or anyone licensing or using the programs on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the Programs, including documentation and technical data, shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement, and to the extent applicable, the additional rights set forth in FAR 52.227-19, Commercial Computer Software—Restricted Rights (June, 1987). Oracle Corporation, 500 Oracle Parkway, Redwood City, CA 94065.

The Programs are not intended for use in any nuclear, aviation, mass transit, medical, or other inherently dangerous applications. It shall be the licensee's responsibility to take all appropriate fail-safe, backup, redundancy, and other measures to ensure the safe use of such applications if the Programs are used for such purposes, and we disclaim liability for any damages caused by such use of the Programs.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

The Programs may provide links to Web sites and access to content, products, and services from third parties. Oracle is not responsible for the availability of, or any content provided on, third-party Web sites. You bear all risks associated with the use of such content. If you choose to purchase any products or services from a third party, the relationship is directly between you and the third party. Oracle is not responsible for: (a) the quality of third-party products or services; or (b) fulfilling any of the terms of the agreement with the third party, including delivery of products or services and warranty obligations related to purchased products or services. Oracle is not responsible for any loss or damage of any sort that you may incur from dealing with

any third party.

• Preface Preface

• Chapter 1 Installing Oracle CODASYL DBMS

- ◆ 1.1 Oracle CODASYL DBMS on OpenVMS I64
- ◆ 1.2 Using Databases from Releases Earlier Than V7.0
- ◆ 1.3 Requirements
- ◆ 1.4 Installation of Oracle CODASYL DBMS Software
- ◆ 1.5 Documentation in Adobe Acrobat Format

• Chapter 2 Enhancements Provided

- ◆ 2.1 Hardware Note on Alpha EV56 Requirement
- ◆ 2.2 Limited IEEE Floating Support in Data Definition Language
- ◆ 2.3 /NOSNAPSHOTS. /[NO]DATA FILE added to DBO/MOVE AREA
- ◆ 2.4 DBO/BACKUP/MULTI/GROUP SIZE Default Value Increased
- ◆ 2.5 The /LOCK TIMEOUT Qualifier Now Allows the Database Default
- ◆ 2.6 Global Statistics Sections for Better Performance
- ◆ 2.7 Maximum Buffer Size Increase
- ◆ 2.8 DBO/SHOW STATISTICS Global Sections Move to P2 Address Space
- ◆ 2.9 EXQUOTA Caused Inaccessible AIJ
- ◆ 2.10 Changes to Default Behavior in Oracle DBMS
- ◆ 2.11 Add option to pass values to /CONFIRM during RESTORE operation.
- ◆ 2.12 Minimum number of allowed users increased

• Chapter 3 Problems Corrected

- ◆ 3.1 DBO/BACKUP/[MULTITHREAD]/ONLINE Bugcheck at DIOBCK\$GET PAGE
- ◆ 3.2 DBO/SHOW STATISTICS Configuration File Issues
- ◆ 3.3 DBO/SHOW STATISTICS Access Violation on Hot Row Information Screen
- ◆ 3.4 Problems If a Full DBO/BACKUP Was Not Done After DBO/MOVE AREA
- ◆ 3.5 DBO Show Statistics Sometimes Bugchecks When Using Process Monitoring
- ◆ 3.6 Incorrect DBO/BACKUP/AFTER Truncate AIJ File Error Handling
- ◆ 3.7 DBO Show Statistics Sometimes Bugchecks on Row Cache Information Screen
- ◆ 3.8 Problem Writing Large TSN Values to Data and Snap Pages
- ◆ 3.9 DBO/RECOVER/ORDER AIJ FILES ERROR HANDLING PROBLEM WHEN ORDERING FILES
- ◆ 3.10 DBO/SHOW STATISTICS/CLUSTER Did Not Show Transaction Type
- ◆ 3.11 DBO Show Statistics not using the full screen size for details
- ◆ 3.12 Node Failure May Cause DBRs to Hang
- ◆ 3.13 DBO/SHOW STATISTICS (Transaction Duration (Total)) playback generates a bugcheck
- ◆ 3.14 %SMG-F-INVROW IN DBO/SHOW STAT USING OPTION WRITE REPORT GRAPH OR BOTH
- ◆ 3.15 Unable to Start DBO/SHOW STATISTICS During Cluster State Transition
- ◆ 3.16 DBO/BACKUP/ONLINE/NOQUIET did not back up the root log file entries
- ◆ 3.17 Unexpected bugcheck dump from DBO/SHOW STATISTICS/HOT STANDBY LOG

• Chapter 4 Known Problems and Restrictions

- ◆ [4.1 Ensure No Processes Have DBMSHR Image Activated](#)
- ◆ [4.2 Syntax for Prestart Transaction Timeout](#)
- ◆ [4.3 Patch Required When Using VMS V8.3 or V8.2-1 on Itanium](#)
- ◆ [4.4 /LOCK_TIMEOUT Documentation Clarification](#)
- ◆ [4.5 Unexpected RCS termination](#)
- ◆ [4.6 Slight Relaxation Of VM\\$MEM RESIDENT USER Requirement](#)
- ◆ [4.7 Patch Required When Using VMS V8.3 and Dedicated CPU Lock Manager](#)
- ◆ [4.8 VM\\$MEM RESIDENT USER Rights Identifier Required](#)
- ◆ [4.9 Features Not Yet Available for OpenVMS I64](#)
- ◆ [4.10 Expect Additional Memory Consumption](#)
- ◆ [4.11 ILINK-E-INVOVRINI Error on I64](#)
- ◆ [4.12 Default Floating Point Format](#)
- ◆ [4.13 SYSTEM-F-IN\\$F\\$MEM Fatal Error With SHARED SYSTEM MEMORY or LARGE MEMORY Enabled in Galaxy Environment](#)
- ◆ [4.14 Oracle CODASYL DBMS and OpenVMS ODS-5 Volumes](#)
- ◆ [4.15 Carryover Locks and NOWAIT Transaction Clarification](#)
- ◆ [4.16 Both Application and Oracle CODASYL DBMS Using SYS\\$HIBER](#)
- ◆ [4.17 Row Cache Not Allowed While Hot Standby Replication is Active](#)
- ◆ [4.18 Exclusive Access Transactions May Deadlock with RCS Process](#)

• [Chapter 5 New Features and Corrections in Previous Releases](#)

Preface

Purpose of This Manual

This manual contains release notes for Oracle CODASYL DBMS release 7.3.1.0. The notes describe changed and enhanced features, upgrade and compatibility information, new and existing software problems and restrictions, and software and documentation corrections.

Intended Audience

This manual is intended for use by all Oracle CODASYL DBMS users. Read this manual before you install, upgrade, or use Oracle CODASYL DBMS release 7.3.1.0.

Document Structure

This manual consists of the following chapters:

<u>Chapter 1</u>	Describes how to install Oracle CODASYL DBMS release 7.3.1.0.
<u>Chapter 2</u>	Describes new and changed features in Oracle CODASYL DBMS release 7.3.1.0.
<u>Chapter 3</u>	Describes problems fixed in Oracle CODASYL DBMS release 7.3.1.0.
<u>Chapter 4</u>	Describes problems, restrictions, and workarounds known to exist in Oracle CODASYL DBMS release 7.3.1.0.
<u>Chapter 5</u>	Describes new features and fixed problems in previous releases.

Conventions

Oracle CODASYL DBMS is often referred to as DBMS in this manual.

HP OpenVMS Industry Standard 64 for Integrity Servers is often referred to as OpenVMS I64.

OpenVMS refers to both OpenVMS Alpha and OpenVMS I64.

Chapter 1

Installing Oracle CODASYL DBMS

All Oracle CODASYL DBMS release 7.3.1.0 kits are full kits. There is no requirement to install any prior release of Oracle CODASYL DBMS prior to installing this release.

1.1 Oracle CODASYL DBMS on OpenVMS I64

In addition to the HP OpenVMS Alpha platform, Oracle CODASYL DBMS is available on the HP OpenVMS Industry Standard 64 for Integrity Servers platform. In general, the Oracle CODASYL DBMS functionality is comparable between the two platforms.

This release provides a full set of Oracle CODASYL DBMS functionality for both platforms, including local and remote database access, as well as native DML and DDL operations. This means that users running on OpenVMS I64 can create Oracle CODASYL DBMS databases, compile, link, and run their database applications natively.

Because the Oracle CODASYL DBMS database format is the same across all supported platforms, you can, for example, back up an Oracle CODASYL DBMS database on an Alpha system, then restore it on an I64 system (the reverse is also true). If necessary, implicit forward conversions are performed to bring the database version to the currently installed level.

With remote access, you can bind to an Oracle CODASYL DBMS database on an Alpha system from an I64 system, or vice versa, as long as the appropriate Oracle CODASYL DBMS software is available on both platforms.

Additionally, if your environment consists of Alpha and I64 systems in a mixed cluster environment, you can access an Oracle CODASYL DBMS release 7.3.1.0 database from either system, or both systems concurrently.

1.2 Using Databases from Releases Earlier Than V7.0

You cannot convert or restore databases from versions earlier than 7.0 directly. The DBO CONVERT command for Oracle CODASYL DBMS 7.3.1.0 supports conversions from V7.0 and V7.1 only.

If you have a V3.3 through V6.1 database, you must convert it to at least V7.0 and then convert it to 7.3.1.0 in two steps. For example, if you have a V4.2 database, install the latest update to DBMS 7.1, convert the database to that version, install DBMS 7.3.1.0 then convert the V7.1 database to 7.3.1.0.

If you attempt to convert or restore a database version prior V7.0 directly to 7.3.1.0, Oracle DBO generates an error.

1.3 Requirements

This version of Oracle CODASYL DBMS supports OpenVMS Alpha 8.4-x and OpenVMS I64 8.4-x.

One of the following conditions must be met in order to install this software:

- OpenVMS Alpha version 8.3 or later
- OpenVMS I64 version 8.3 or later.

1.4 Installation of Oracle CODASYL DBMS Software

Please refer to the *CODASYL DBMS 7.3 Installation Guide* for detailed Oracle CODASYL DBMS installation instructions. Oracle strongly recommends that you read the installation guide before attempting an installation.

To extract either the PostScript (PS) or text (TXT) version of the installation guide from the kit, use one of the following commands:

For OpenVMS Alpha:

```
$ BACKUP <device>:DBM07310A073.A/SAVE/SEL=DBM073_INSTALL_GDE.PS  
$ BACKUP <device>:DBM07310A073.A/SAVE/SEL=DBM073_INSTALL_GDE.TXT
```

For OpenVMS I64:

```
$ BACKUP <device>:DBM07310I073.A/SAVE/SEL=DBM073_INSTALL_GDE.PS  
$ BACKUP <device>:DBM07310I073.A/SAVE/SEL=DBM073_INSTALL_GDE.TXT
```

The release 7.3 installation guide is available on MetaLink and OTN in Adobe Acrobat PDF format.

1.5 Documentation in Adobe Acrobat Format

You can view the documentation in Adobe Acrobat format using the Acrobat Reader, which allows anyone to view, navigate, and print documents in the Adobe Portable Document Format (PDF). For information about obtaining a free copy of Acrobat Reader and for information on supported platforms, see the Adobe Web site at:

<http://www.adobe.com>

The Oracle CODASYL DBMS and Hot Standby documentation in Adobe Acrobat format is available on MetaLink and OTN.

Chapter 2

Enhancements Provided

This chapter describes new and changed features in Oracle CODASYL DBMS release 7.3.1.0.

2.1 Hardware Note on Alpha EV56 Requirement

Oracle CODASYL DBMS Release 7.3.1.0 has been optimized for the Alpha EV56 platform as a minimum. If running on an older platform, Oracle CODASYL DBMS will still execute properly, but will run some operations in emulation mode and performance will be significantly slower. Oracle supports running Oracle DBMS in this emulation mode as long as any problems encountered can be reproduced on an EV56 or greater system.

2.2 Limited IEEE Floating Support in Data Definition Language

Starting in 7.3.1.0, Oracle CODASYL DBMS provides support for IEEE single and double precision floating point data types within the Data Definition Language (DDL).

Two new keywords have been added to the DDL TYPE clause:

- TYPE IS S_FLOATING ! single-precision floating point
- TYPE IS T_FLOATING ! double-precision floating point

Example:

```
RECORD IS FOO0677
  CHECK IS FOO0677_A EQ FOO0677_B
  ITEM FOO0677_A
    TYPE IS SIGNED BYTE
  ITEM FOO0677_B
    TYPE IS S_FLOATING

RECORD IS FOO0705
  CHECK IS FOO0705_A EQ FOO0705_B
  ITEM FOO0705_A
    TYPE IS PACKED DECIMAL 10 -2
  ITEM FOO0705_B
    TYPE IS T_FLOATING
```

The new data type can be used in SCHEMA and STORAGE_SCHEMA declarations. They can also be used within SUBSCHEMA declarations if the target host language supports the datatype. If an IEEE data item is specified in the SCHEMA or STORAGE_SCHEMA, but is not supported by the host language, you will need to specify a supported comparable data type in the SUBSCHEMA, allowing DBMS to convert internally between the two.

2.3 /NOSNAPSHOTS, /[NO]DATA_FILE added to DBO/MOVE_AREA

Currently DBO/MOVE_AREA moves or creates a new version of BOTH the storage area data (*.DBS) and snapshot (*.SNP) files. This new syntax allows moving ONLY the data area file or ONLY the snapshot area file for all or for named storage areas. /NODATA_FILE and /NOSNAPSHOTS are positional qualifiers that can be specified globally as a default and/or for one or more named storage areas. They can be specified on the command line or in an options file using the existing DBO/MOVE_AREA /OPTION=filespec qualifier.

The syntax for these qualifiers is as follows –

```
/[NO]SNAPSHOTS[=( [FILE=filespec][ALLOCATION=n] )]
```

NOSNAPSHOTS does not move the storage area snapshot file(s). It only moves the data storage area file(s). SNAPSHOTS is the default.

[=([FILE=filespec][ALLOCATION=n])] cannot be specified with NOSNAPSHOTS.

SNAPSHOTS[=([FILE=filespec][ALLOCATION=n])] is an existing qualifier but now it can be negated. SNAPSHOTS as a local qualifier can override NOSNAPSHOTS as a global qualifier. NOSNAPSHOTS as a local qualifier can override SNAPSHOTS as a global qualifier.

```
/[NO]DATA_FILE
```

NODATA_FILE does not move the storage area data file(s). It only moves the snapshot storage area file(s). DATA_FILE is the default. It does not accept any values. DATA_FILE as a local qualifier can override NODATA_FILE as a global qualifier. NODATA_FILE as a local qualifier can override DATA_FILE as a global qualifier.

If NODATA_FILE is specified the storage area data file is not moved or modified. A new version of the file will not be created. If NOSNAPSHOT is specified the storage area snapshot file is not moved or modified. A new version of the file will not be created. Any existing DBO/MOVE_AREA qualifiers that would require an update/change to the data area file are disallowed if /NODATA_FILE is specified and any qualifiers that would require an update/change to the snapshot area file are disallowed if /NOSNAPSHOTS is specified.

Therefore the following existing /MOVE_AREA qualifiers cannot be specified with either /NODATA_FILE or /NOSNAPSHOTS:

- /ROOT
- /BLOCKS_PER_PAGE
- /NODES_MAX
- /USERS_MAX

The following existing /MOVE_AREA qualifiers cannot be specified with /NODATA_FILE:

- /FILE
- /SPAMS
- /THRESHOLDS
- /EXTENSION

Oracle® CODASYL DBMS for OpenVMS

The following existing /MOVE_AREA qualifiers cannot be specified with /NOSNAPSHOTS:
/SNAPSHOTS=(FILE=filespec) /SNAPSHOTS=(ALLOCATION=n)

In the following example only the storage area snapshot files are moved for all database storage areas.

```
$ DBO/move_area/all/nodata_file/nolog/dir=[.move] parts.ROO
%DBO-W-DOFULLBCK, full database backup should be done to ensure future recovery
```

In the following example only the storage area data files are moved for all database storage areas.

```
$ DBO/move_area/all/nosnapshots/nolog/dir=[.move] parts.ROO
%DBO-W-DOFULLBCK, full database backup should be done to ensure future recovery
```

In the following example only the snapshot storage area file is moved for the EMPLOYEE storage area and only the data storage area file is moved for the MARKET storage area. Note that for the MARKET storage area /DATA_FILE did not need to be specified since it is the default.

```
$ DBO/move_area/nolog parts.ROO -
  EMPLOYEE /nodata_file -
            /snapshots=(file=DISK:[DIRECTORY]test_employee.snp), -
  MARKET  /data_file -
            /file=DISK:[DIRECTORY]test_market -
            /nosnapshots
%DBO-W-DOFULLBCK, full database backup should be done to ensure future recovery
```

In the following example an options file is used to specify the storage areas to be moved. Only the data storage area file is moved for EMPLOYEE, only the snapshot storage area file is moved for MARKET, and both the snapshot and data storage area files are moved for DEPARTMENTS. NOTE that /DATA_FILE and /SNAPSHOT are the defaults.

```
$ DBO/move_area/nolog/directory=DISK:[DIRECTORY]/option=testmove.opt -
parts.ROO
  EMPLOYEE -
            /file=DISK:[DIRECTORY]test_employee.dbs -
            /nosnapshot -
  MARKET  /nodata_file -
            /snapshot = (file=DISK:[DIRECTORY]test_market.snp)
  DEPARTMENTS -
            /file=DISK:[DIRECTORY]test_departments -
            /snapshot = (file=DISK:[DIRECTORY]test_departments.snp)
%DBO-W-DOFULLBCK, full database backup should be done to ensure future recovery
```

In the following example the global default qualifiers designate that only the snapshot files should be moved for all storage areas. However, an options file is used to override the default for specific storage areas. Therefore, only the data storage area file is moved for EMPLOYEE, only the snapshot storage area file is moved for MARKET, and both the snapshot and data storage area files are moved for DEPARTMENTS. NOTE that in this case /DATA_FILE needed to be specified in the options file to override the global specification of /NODATA_FILE but /NODATA_FILE did not have to be specified in the options file. Also /NOSNAPSHOT had to be specified in the options file to override the assumed global default of /SNAPSHOT.

```
$ DBO/move_area/all/directory=DISK:[DIRECTORY]/nolog/nodata_file-
/option=testmove.opt parts.ROO
  EMPLOYEE /data_file -
            /file=DISK:[DIRECTORY]test_employee.dbs -
            /nosnapshot
```

Oracle® CODASYL DBMS for OpenVMS

```
MARKET      /nodata_file -
            /snapshot = (file=DISK:[DIRECTORY]test_market.snp)
DEPARTMENTS -
            /data_file -
            /file=DISK:[DIRECTORY]test_departments -
            /snapshot = (file=DISK:[DIRECTORY]test_departments.snp)
%DBO-W-DOFULLBCK, full database backup should be done to ensure future recovery
```

2.4 DBO/BACKUP/MULTI/GROUP_SIZE Default Value Increased

The DBO/BACKUP/MULTI/GROUP_SIZE qualifier specifies the frequency at which XOR recovery blocks are written to tape. In prior releases of Oracle CODAYSL DBMS, the default group size when writing to a tape device was 10.

Starting with Oracle DBMS Release 7.3.1.0, the default group size when writing to a tape device is 100. This change may result in an increased throughput and a reduction in tape space used.

2.5 The /LOCK_TIMEOUT Qualifier Now Allows the Database Default

For the Oracle DBMS DBO commands DBO/BACKUP/MULTI/ONLINE, DBO/COPY/ONLINE, the /LOCK_TIMEOUT qualifier can be specified. Previously the /LOCK_TIMEOUT qualifier required a value, the maximum time in seconds to wait for acquiring the database QUIET POINT and other locks used for online database access. If "/LOCK_TIMEOUT = #" was not specified, DBO would wait indefinitely to acquire the database lock it needed.

Now the /LOCK_TIMEOUT qualifier can be specified without a value. In this case the default lock timeout value specified for the database will be used. Specifically, the default lock timeout value used will be the value of the logical name DBM\$BIND_LOCK_TIMEOUT_INTERVAL if it has been specified. If DBM\$BIND_LOCK_TIMEOUT_INTERVAL has not been specified, value used will be the maximum possible lock timeout value which can be specified for an Oracle DBMS database.

The new syntax for this qualifier is as follows –

```
/LOCK_TIMEOUT [ = # ]
```

Note that /LOCK_TIMEOUT is not the default and must be specified. The default if /LOCK_TIMEOUT is not specified continues to be to wait indefinitely to acquire the QUIET POINT or other database locks requested by DBO.

2.6 Global Statistics Sections for Better Performance

On systems with many CPUs updating database statistics from many application processes causes memory cache invalidation and therefore prolongs the update of the statistics data.

With this change the DBMS Monitor creates 16 global statistic sections for systems with 16 or more CPUs. Application processes attach to a statistics section based on the modulo 16 of their process ID value.

This should reduce the coincidence that two or more processes use the same global section from different processors and thus causing memory cache invalidation when updating statistics data.

The default used for RAD (Resource Allocation Domain) systems still remains (see below).

The use of multiple global statistic sections can be overridden with the following system logical name:

```
$ DEFINE /SYSTEM DBM$BIND_MONITOR_GLOBAL_STATS_SECTIONS n
  n = 0 - always use statistics in the database's shared memory section
  n = 1..16 - use statistics in separate global sections
              with n the number of global sections being used
  If n is equal -1 or if the logical is not defined
  use the default (see below).
```

By default the statistics area in the database's shared memory section is used unless a system has more than one RAD with memory or, the system has 16 or more CPUs.

In the case of more than one RAD with memory one global statistics section is created per RAD with memory. In the case of 16 or more CPUs 16 global statistics sections are created. The more than one RAD with memory case has precedence over the 16 or more CPUs case.

2.7 Maximum Buffer Size Increase

Previously, the maximum allowed database buffer size was 128 blocks. This limit has been increased. The current maximum allowed database buffer size is 256 disk blocks.

Be aware that using larger database buffer sizes will require additional virtual memory.

2.8 DBO/SHOW STATISTICS Global Sections Move to P2 Address Space

In releases prior to 7.3, the global sections created by DBO/SHOW/STATISTICS were mapped to P0 space. Starting with this release, those global sections will automatically be mapped to P2 space.

This change should help alleviate any VASFULL errors that may have previously occurred.

2.9 EXQUOTA Caused Inaccessible AIJ

In prior versions of Oracle DBMS, if an After Image Journal (AIJ) write operation failed with a FILACCERR error and COMMIT TO JOURNAL was enabled, the AIJ file would be immediately marked as inaccessible, the database would be shutdown, and the offending process (database user or database server) would be terminated. Manual intervention would be required to reset journalling before the database could be re-opened.

Starting in this release, Oracle DBMS has eased one of the restrictions. Now, if a user process gets an FILACCERR error due to "exceeded quota" (EXQUOTA), the AIJ will not be marked as inaccessible and the database won't be shutdown. The process will, however, still be terminated. This abnormal termination will cause a Database Recovery (DBR) process to be automatically started to recover that user. No manual intervention would be required and normal database processing would continue after the DBR completes.

Note that this change only will affect database user processes. The behavior of Database server processes (such as ALS, RCS, LCS, LRS) remains the same.

2.10 Changes to Default Behavior in Oracle DBMS

This release of Oracle DBMS changes the following default behavior.

DBO/CREATE and DBO/MODIFY

These new defaults will be used when creating a database.

- The default PAGE SIZE changes from 2 to 4 blocks
- The default BUFFER SIZE changes from 3 pages to 4 pages
- The default NUMBER OF BUFFERS changes from 20 to 250 buffers
- The default NUMBER OF RECOVERY BUFFERS changes from 20 to 250 buffers

These new defaults do not affect databases created in DBMS V7.2 (or older versions) and converted to CODASYL Oracle DBMS 7.3.1.0 using DBO/CONVERT or DBO/RESTORE/MULTI.

2.11 Add option to pass values to /CONFIRM during RESTORE operation.

In prior releases of Oracle CODASYL DBMS, if problems occur during tape restore operations and the /CONFIRM qualifier is specified, DBO will wait for user input on the terminal.

```
$ DBO/restore/multi/multi/nocdd/rewind lma1001:vol002.DBF PARTS /directory=[] /log
%DBO-I-WRNGLBL, Tape on LMA1001 was incorrectly labeled. Expected VOL002 - Found MF_PER
%DBO-I-TAPEDISPR, Specify tape disposition for LMA1001 (QUIT,OVERRIDE,RETRY,UNLOAD)
DBO> QUIT (User has to enter the RESPONSE.)
%DBO-F-ABORT, operator requested abort on fatal error
%DBO-F-FATALERR, fatal error on RESTORE
%DBO-F-FTL_RSTR, Fatal error for RESTORE operation at 13-JUL-2009 11:22:32.90
```

```
$ DBO/restore/multi/multi/nocdd/rewind lma1001:vol002.DBF PARTS /directory=[] /log
%DBO-I-WRNGLBL, Tape on LMA1001 was incorrectly labeled. Expected VOL002 - Found MF_PER
%DBO-I-TAPEDISPR, Specify tape disposition for LMA1001 (QUIT,OVERRIDE,RETRY,UNLOAD)
DBO> RETRY (User has to enter the RESPONSE.)
%MOUNT-I-MOUNTED, MF_PER mounted on LMA1001:
%DBO-I-WRNGLBL, Tape on LMA1001 was incorrectly labeled. Expected VOL002 - Found MF_PER
%DBO-I-TAPEDISPR, Specify tape disposition for LMA1001 (QUIT,OVERRIDE,RETRY,UNLOAD)
DBO> QUIT (User has to enter the RESPONSE.)
%DBO-F-ABORT, operator requested abort on fatal error
%DBO-F-FATALERR, fatal error on RESTORE
%DBO-F-FTL_RSTR, Fatal error for RESTORE operation at 13-JUL-2009 11:22:55.86
```

```
$ DBO/restore/multi/multi/nocdd/rewind lma1001:vol002.DBF PARTS /directory=[] /log
%DBO-I-WRNGLBL, Tape on LMA1001 was incorrectly labeled. Expected VOL002 - Found MF_PER
%DBO-I-TAPEDISPR, Specify tape disposition for LMA1001 (QUIT,OVERRIDE,RETRY,UNLOAD)
DBO> OVERRIDE (User has to enter the RESPONSE.)
%DBO-F-FILACCERR, error opening input file LMA1001:[000000]VOL002.DBF;
-SYSTEM-W-NOSUCHFILE, no such file
%DBO-F-FATALERR, fatal error on RESTORE
%DBO-F-FTL_RSTR, Fatal error for RESTORE operation at 13-JUL-2009 11:23:05.59
```

A new feature has been added in this release to correct this problem. The user has the option of selecting values for /CONFIRM during a RESTORE from tape operation. The new syntax and valid values are:

```
DBO/RESTORE/MULTI/MULTI... /CONFIRM[=QUIT|RETRY=x|OVERRIDE|UNLOAD]
```

```
$ DBO/restore/multi/multi/nocdd/rewind lma1001:vol002.DBF PARTS /directory=[] /log /confirm=quit
%MOUNT-I-MOUNTED, MF_PER mounted on LMA1001:
%DBO-I-TAPEDEF, Terminating restore operation as requested by user
%DBO-F-ABORT, operator requested abort on fatal error
%DBO-F-FATALERR, fatal error on RESTORE
%DBO-F-FTL_RSTR, Fatal error for RESTORE operation at 13-JUL-2009 11:43:31.35
```

```
$ DBO/restore/multi/multi/nocdd/rewind lma1001:vol002.DBF PARTS /directory=[] /log /confirm=retr
%DBO-I-TAPEDEF, Retrying tape operation as requested by user
%MOUNT-I-MOUNTED, MF_PER mounted on LMA1001:
%DBO-I-TAPEDEF, Retrying tape operation as requested by user
%MOUNT-I-MOUNTED, MF_PER mounted on LMA1001:
%DBO-F-ABORT, operator requested abort on fatal error
%DBO-F-FATALERR, fatal error on RESTORE
%DBO-F-FTL_RSTR, Fatal error for RESTORE operation at 13-JUL-2009 11:43:42.97
```

```
$ DBO/restore/multi/multi/nocdd/rewind lma1001:vol002.DBF PARTS /directory=[]V/log /confirm=over
```

Oracle® CODASYL DBMS for OpenVMS

```
%DBO-I-TAPEDEF, Overriding tape label as requested by user  
%DBO-F-FILACCERR, error opening input file LMA1001:[000000]VOL002.DBF;  
-SYSTEM-W-NOSUCHFILE, no such file %DBO-F-FATALERR, fatal error on RESTORE  
%DBO-F-FTL_RSTR, Fatal error for RESTORE operation at 13-JUL-2009 11:43:58.38
```

2.12 Minimum number of allowed users increased

In prior versions of Oracle CODASYL DBMS, the minimum number of allowed database users was one. This minimum has been increased to five to allow for various optional database servers (such as the ABS or RCS or ALS) to access the database.

Chapter 3

Problems Corrected

This chapter describes software errors corrected in Oracle CODASYL DBMS release 7.3.1.0.

3.1 DBO/BACKUP[/MULTITHREAD]/ONLINE Bugcheck at DIOBCK\$GET_PAGE

Bug 12741822

Starting with Oracle CODASYL DBMS Release 7.2.4.2, it was possible for an online database backup to return an error with a bugcheck exception at DIOBCK\$GET_PAGE.

The error indicates that there was insufficient free space on the data page being synthesized for backup to hold the line and TSN indices for that page. This was caused by incorrectly including the size of freed indices in the calculation of required bytes.

The problem only occurs when using the /ONLINE qualifier on the backup command. Although it is difficult to predicate if a backup failure will occur, those data pages with little free space and a higher rate of deleted records may be more susceptible to this error. Restarting the backup may be successful or may fail again on the same or different data page.

This problem has been corrected in Oracle CODASYL DBMS release 7.3.1.0.

3.2 DBO/SHOW STATISTICS Configuration File Issues

12710800 and 12710931

Starting in release V7.2-50, attempting to use a Configuration file on the DBO/SHOW STATISTICS command would generate an access violation. The workaround was to remove that Configuration file specification and use the equivalent command line qualifiers where possible.

This problem has been corrected in Oracle CODASYL DBMS release 7.3.1.0.

3.3 DBO/SHOW STATISTICS Access Violation on Hot Row Information Screen

Starting in release V7.2–50, attempting to view "the Hot Row Information" screen will cause an access violation. This problem does not affect other DBO/SHOW STATISTICS functions. There is no workaround for this problem.

This problem has been corrected in Oracle CODASYL DBMS release 7.3.1.0.

3.4 Problems If a Full DBO/BACKUP Was Not Done After DBO/MOVE_AREA

BUG 13070493

Oracle recommends taking a full database backup after a DBO/MOVE operation. Failure to do so may cause data corruption.

For example, suppose a full database backup was taken prior to the DBO/MOVE and an incremental database backup was taken after the DBO/MOVE. If the full backup was retored, followed by the incremental restore, changes made to the root file by DBO/MOVE would be lost and the database could be left in a corrupt state.

If the full backup had been taken after the DBO/MOVE operation, the problem would not have occurred.

This problem has now been fixed in Oracle CODASYL DBMS release 7.3.1.0. The DBO/MOVE command will now warn you to take a full backup at completion. Also, attempting an incremental backup after the DBO/MOVE, but before taking a full backup will not be allowed.

```
$ DBO/MOVE_AREA/DIRECTORY=[.move_db]/NOLOG
  device:[directory]PARTS.ROO PARTS
%DBO-W-DOFULLBCK, full database backup should be done to ensure future recovery
%DBO-I-COMPLETED, MOVE_AREA operation completed at 11-OCT-2011 12:09:15.40
$
$ DBO/backup/multi/online/nolog/incremental PARTS PARTS_INC
%DBO-F-NOFULLBCK, no full backup of this database exists
%DBO-F-FTL_BCK, Fatal error for BACKUP operation at 11-OCT-2011 12:30:18.08
```

3.5 DBO Show Statistics Sometimes Bugchecks When Using Process Monitoring

Bug 13070947

The Oracle CODASYL DBMS DBO Show Statistics command "Process monitoring" option would sometimes bugcheck. The following sequence of commands while in an DBO/SHOW STATISTICS window would sometimes result in a bugcheck.

```
! (to get to the Select Tool)
R. <<more>>
J. Process monitoring
B. Activate specific eligible process
Enter process ID to "activate" (or "?"): (enter id here)

! (to get to the Select Tool)
J. Process monitoring
A. Select activated process to monitor
Process detached; Continue reviewing? ("Yes" to continue, "RETURN" to cancel)
```

The bugcheck looked like the following:

```
***** Exception at FFFFFFFF84230400 : Image LIBOTS + 00002400
%SYSTEM-F-ACCVIO, access violation, reason mask=00, virtual address=0000000003A64000, PC=FFFFFFFF
PS=0000001B
Saved PC = 000000008089AB10 : DBO72\KUTDIS$TOOLS_CONTROL + 0000EE80
Saved PC = 00000000809855A0 : DBO72\KUTDIS$TOOLS_MENU + 0000DBC0
```

This problem has been corrected in Oracle CODASYL DBMS Release 7.3.1.0. DBO will no longer bugcheck and will show the requested statistics.

3.7 DBO Show Statistics Sometimes Bugchecks on Row Cache Information Screen

Bug 13558012

The Oracle CODASYL DBMS DBO Show Statistics command "Hot Row Information" option would sometimes bugcheck. The following sequence of commands while in an DBO/SHOW STATISTICS window would sometimes result in a bugcheck when monitoring row cache.

```
M (Menu)
X (Row Cache Information)
B (Hot Row Information)
Select one cache
```

The bugcheck looked like the following:

```
SYSTEM-F-ACCVIO, access violation, virtual address=FFFFFFFF817B80B0
Exception occurred at DBO72\KUTDIS$SETUP_RCD_ACTRTN + 00000460
```

This problem has been corrected in Oracle CODASYL DBMS Release 7.3.1.0. DBO will no longer bugcheck and will show the requested statistics.

3.8 Problem Writing Large TSN Values to Data and Snap Pages

\BUG12741822)

Bug 12741822

A problem has been discovered that could affect how Transaction Sequence Numbers (TSNs) are stored on data and snapshot pages. This problem only occurs when TSNs have a value greater than 2,147,483,647.

A TSN is a sequentially increasing number associated with an Rdb database transaction. A unique TSN is assigned at the start of an update transaction and represents when the transaction started relative to other transactions.

TSNs are used throughout DBMS memory data structures. Additionally, DBMS writes the TSN to on-disk structures such as data files, snapshot files, and after-image journal (AIJ) files to reflect which transaction modified a particular row. Read-only transactions use the TSNs stored on data and snapshot pages to determine which version of a modified row can be seen by the reader.

The problem occurs due to a change to DBMS Version 7.0 in how TSNs are stored on data and snap pages. Prior to that version, TSNs were 32-bit integers with a maximum value of 4,294,967,295. With 7.0, the size of the TSN field was changed to be a 64-bit integer. So as not to require a database unload and reload to support the larger field size, an algorithm was devised to break up those TSNs that would no longer fit into a 32-bit field into 2 pieces when the page is updated. The original value would be regenerated prior to reading from the page.

Recently, it was discovered that this algorithm could cause such large TSNs to possibly be stored incorrectly. Occurrences would be more prevalent during the window where some active transaction would have TSN values greater than the 2,147,483,647 value, while others had values less than that maximum.

The problem may manifest itself with errors during an DBO /VERIFY /ONLINE, or with incorrect data returned from read-only transactions. Additionally, a DBO /BACKUP [/MULTITHREAD] /ONLINE may produce a backup file that will restore a corrupted database (although an DBO /RECOVER of all spanning after-journal files will usually fix the corruption).

If you believe that your database exhibits these symptoms, you can use the DBO /INIT /TSN command to reset the TSNs. As always, Oracle recommends taking a full off-line database backup prior to any rootfile modifications.

Note

A Support article about this issue and how to resolve it has been written also. Please reference document number 1440100.1 on My Oracle Support to read this article. Support provides a command procedure that can patch the TSNBLK and CSNBLK numbers to get around the problem.

The problem has been corrected in Oracle CODASYL DBMS Release 7.3.1.0. TSNs will now be stored correctly on data pages.

3.9 DBO/RECOVER/ORDER_AIJ_FILES ERROR HANDLING PROBLEM WHEN ORDERING FILES

There was an error handling problem if an error occurred when DBO/RECOVER/ORDER_AIJ_FILES was sorting the input AIJ files in ascending order by sequence number. An access violation occurred in the error handler followed by a loop producing multiple access violations which produced multiple DBOBUGCHK.DMP files. A "CTRL/Y" had to be executed to terminate the loop. This problem has now been fixed.

The following example shows the problem. An error exception occurs when DBO/RECOVER is sorting the input AIJ files in ascending order by sequence number at the start of the recovery operation because one of the AIJ files is locked by another user. The error handler for this error exception gets into a loop which produces multiple "%SYSTEM-F-ACCVIO" errors and multiple DBOBUGCHK.DMP files. This loop can only be terminated by a "CTRL/Y".

```
$ DBO/RECOVER AIJB* /ORDER_AIJ_FILES/NOLOG
%DBO-F-FILACCERR, error reading journal file
DEVICE:[DIRECTORY]AIJB_4.AIJ;1
-RMS-E-FLK, file currently locked by another user
%DBO-F-FTL_RCV, Fatal error for RECOVER operation at 13-FEB-2012
13:37:29.91
%SYSTEM-F-ACCVIO, access violation, reason mask=00, virtual
address=000000000000018C, PC=000000008078F4A1, PS=0000001B
%DBO-F-FATALOSI, Fatal error from the Operating System Interface.
%DBO-I-BUGCHKDMP, generating bugcheck dump file
DEVICE:[DIRECTORY]DBOBUGCHK.DMP;
%DBO-F-FTL_RCV, Fatal error for RECOVER operation at 13-FEB-2012
13:37:30.03
%SYSTEM-F-ACCVIO, access violation, reason mask=00, virtual
address=000000000000018C, PC=000000008078F4A1, PS=0000001B
%DBO-F-FATALOSI, Fatal error from the Operating System Interface.
%DBO-I-BUGCHKDMP, generating bugcheck dump file
DEVICE:[DIRECTORY]DBOBUGCHK.DMP;
%DBO-F-FTL_RCV, Fatal error for RECOVER operation at 13-FEB-2012
13:37:30.11
%SYSTEM-F-ACCVIO, access violation, reason mask=00, virtual
address=000000000000018C, PC=000000008078F4A1, PS=0000001B
%DBO-F-FATALOSI, Fatal error from the Operating System Interface.
%DBO-I-BUGCHKDMP, generating bugcheck dump file
DEVICE:[DIRECTORY]DBOBUGCHK.DMP;
%DBO-F-FTL_RCV, Fatal error for RECOVER operation at 13-FEB-2012
13:37:30.19
%SYSTEM-F-ACCVIO, access violation, reason mask=00, virtual
address=000000000000018C, PC=000000008078F4A1, PS=0000001B
%DBO-F-FATALOSI, Fatal error from the Operating System Interface.
%DBO-I-BUGCHKDMP, generating bugcheck dump file
DEVICE:[DIRECTORY]DBOBUGCHK.DMP_MINI;
%DBO-F-FTL_RCV, Fatal error for RECOVER operation at 13-FEB-2012
13:37:30.25
```

The following example shows that this problem has been fixed. Now when the AIJ file to be sorted is locked by another user the error exception is handled correctly. The fatal error is output and the recovery is aborted.

```
$ DBO/RECOVER AIJB* /ORDER_AIJ_FILES/NOLOG
%DBO-F-FILACCERR, error reading journal file
DEVICE:[DIRECTORY]AIJB_4.AIJ;1
```

Oracle® CODASYL DBMS for OpenVMS

-RMS-E-FLK, file currently locked by another user
%DBO-F-FTL_RCV, Fatal error for RECOVER operation at 13-FEB-2012
13:40:40.26

This problem has been corrected in Oracle CODASYL DBMS release 7.3.1.0.

3.10 DBO/SHOW STATISTICS/CLUSTER Did Not Show Transaction Type

Bug 10009143

In releases prior to Oracle CODASYL DBMS Release 7.3.1.0, there was a problem in the DBO/SHOW STATISTICS/CLUSTER command where the Transaction Type was not being displayed on the Checkpoint Information screen, when the transaction was executed on a node other than the current node. This problem only occurred when the DBO command was executed on an Itanium system.

This problem has been corrected in Oracle CODASYL DBMS Release 7.3.1.0. DBO now shows the transaction type on the Checkpoint Information screen.

3.11 DBO Show Statistics not using the full screen size for details

Bug 13892668, 13359835

In prior versions of Oracle CODASYL DBMS, the DBO Show Statistics display did not fully use the size of the display, and in some cases limited the output to 17 lines of details.

This problem has been corrected in Oracle CODASYL DBMS Release 7.3.1.0. DBO Show Statistics now uses more screen space when available. The screen size is determined by the terminal page size (SET TERMINAL/PAGE) or when using the /ROW qualifier on DBO Show Statistics.

3.12 Node Failure May Cause DBRs to Hang

Bug 14582053

Under normal conditions, if Oracle CODASYL DBMS is running in a cluster environment and one of the members fails, database recovery processes (DBR) will automatically be created on one of the other cluster members, where the database is currently open, to recover the database users running on that failed node.

In rare circumstances, it was possible for one or more of those DBR's to become blocked waiting for a database page lock held by an active user. That user would be blocked waiting for cluster membership (PR). This scenario would be more likely occur when the systems were under a heavy load, such that the recovering node is resource constrained and creating recovery processes in such an environment takes longer than usual.

The workaround would be to DELPRC or STOP/ID all those blocker processes, or stop all database activity on the remaining nodes using DBO/CLOSE/CLUSTER /ABORT=DELPRC <database-name>, then reopen the database. The next attach would cause DBRs to be created and recovery should complete successfully.

This problem has been corrected in Oracle CODASYL DBMS Release 7.3.1.0. The blocker transaction will now release the contested resource, allowing the recovery to continue.

3.13 DBO/SHOW STATISTICS (Transaction Duration (Total)) playback generates a bugcheck

Bug 14549269

In prior versions of Oracle CODASYL DBMS the DBO Show Statistics (Transaction Duration (Total)) screen on Integrity systems generates bugcheck with an arithmetic trap

```
%SYSTEM-F-INTDIV, arithmetic trap, integer divide by zero at  
PC=FFFFFFFF80A91E51, PS=0000001B  
%DBO-F-FATALOSI, Fatal error from the Operating System Interface.  
%DBO-I-BUGCHKDMP, generating bugcheck dump file DISK1:[TESTING]DBOBUGCHK.DMP;  
%DBO-F-FTL_SHOW, Fatal error for SHOW operation at 24-AUG-2012 06:29:45.45
```

This problem has been corrected in Oracle CODASYL DBMS Version 7.3.1.0. DBO Show Statistics no longer generates a fatal error for this screen.

3.14 %SMG-F-INVROW IN DBO/SHOW STAT USING OPTION WRITE REPORT GRAPH OR BOTH

Bug 14807019

When using DBO/SHOW STATISTICS on a database (e.g. PARTS) and doing M (Menu), then a letter that gives a screen in which you can choose O (Option) to write a Report with Graphs or Numbers or Both, and you choose A (graph) or C (Both) the report is written partially and then \$DBO/SHOW STATISTICS crashes with the error message:

```
%SMG-F-INVROW, invalid row
%DBO-F-FATALOSI, Fatal error from the Operating System Interface.
%DBO-F-FTL_SHOW, Fatal error for SHOW operation <<DATE & TIME >>
```

This problem has been corrected in Oracle CODASYL DBMS Version 7.3.1.0. DBO Show Statistics no longer generates an invalid row fatal error for this screen.

3.15 Unable to Start DBO/SHOW STATISTICS During Cluster State Transition

In prior versions of Oracle CODASYL DBMS, an attempt to run DBO/SHOW STATISTICS during a cluster state transition would cause the process to hang until the transition was complete.

BUG 14582053

Under normal circumstances, if an DBMS database, opened on multiple cluster members, was abnormally closed on one node with active users, a Database Recover process (DBR) would be initiated on another node to recover those failed users. The abnormal closure could be caused by a number of situations including: node failure, DBMS monitor failure, or an DBO/CLOSE/ABORT=DELPRC command.

During this recovery period, all database processes are frozen and new attaches (including DBO/SHOW STATISTICS) are stalled. When the recovery completes, the database has normalized and activity can continue on the remaining available nodes (cluster state transition).

In very rare circumstances, it may be possible for the DBR process to become stalled, typically on an undetected deadlock. If the DBO SHOW/STATISTICS utility had been able to start, the "Stall Messages" and "DBR Activity" screens could help determine the nature of the stall, and the Tools facility could be used terminate the processes blocking the DBR.

This problem has now been fixed. DBO/SHOW STATISTICS is no longer subject to the freeze protocol and thus can start and collect information even when there is a stalled recovery process.

3.16 DBO/BACKUP/ONLINE/NOQUIET did not back up the root log file entries

The LOGFIL structure in the Oracle CODASYL DBMS database root file contains entries for different DBMS server log files. These entries can be initialized by the DBO/SET SERVER COMMAND. Due to a problem with an DBO/BACKUP[/MULTITHREAD]/ONLINE/NOQUIET backup, the log file entries in the root were not backed up to the *.DBF backup file. When the database was restored the root log file entries were not corrupt but were missing and had to be redefined. This only happened if the /NOQUIET qualifier was specified for the backup.

This problem has been fixed in Oracle CODASYL DBMS Release 7.3.1.0. Now DBO/BACKUP/ONLINE/NOQUIET correctly backs up the LOGFIL log file entries in the database root so that they are all restored correctly.

3.17 Unexpected bugcheck dump from DBO/SHOW STATISTICS/HOT_STANDBY_LOG

Bug 15883785

In prior versions of Oracle CODASYL DBMS, the DBO Show Statistics command with /Hot_Standby_Log qualifier may generate an access violation (ACCVIO) and produce a bugcheck dump. The bugcheck summary would be similar to this example:

- Itanium OpenVMS 8.4
- Oracle CODASYL DBMS Server 7.2.5.2.0
- Got a DBOBUGCHK.DMP
- SYSTEM-F-ACCVIO, access violation, virtual address=0000000050435400
- Exception occurred at symbol not found
- Called from DBO72\KUTDIS\$HS_NOTIFY + 00000840
- Called from DBO72\KUTDIS\$EVENT_NOTIFY + 00000220
- Called from DBO72\KUTDIS\$DISPLAY_ASTX + 00000D00
- Running image DBO72.EXE
- Command line was: DBO/SHOW STATI/TIME=5/NOINTERACTIVE-
- /UNTIL=14-NOV-2012 16:59:35.42/HOT_STANDBY=SYS\$LOGIN:GET_HS.LOG
TEST_DATABASE

This problem has been corrected in Oracle CODASYL DBMS Release 7.3.1.0.

Chapter 4

Known Problems and Restrictions

This chapter describes problems and restrictions relating to Oracle COADASYL DBMS release 7.3.1.0 and includes workarounds where appropriate.

4.1 Ensure No Processes Have DBMSHR Image Activated

The Oracle CODASYL DBMS installation procedure checks to make sure that the DBMS Monitor (DBMMON) process is not running. However, it is also important to make sure that there are no processes on the cluster that have images activated with a prior version DBMSHR image. Such processes may not be currently attached to a database but may do so in the future and could cause problems by using an older DBMSHR image with a later DBMS installation.

The following command procedure can be used on each cluster node that shares the system disk to determine if there are any processes that have activated the DBMSHR72.EXE image. This procedure should be executed by a privileged account after MONSTOP72.COM has been run. Any processes that have DBMSHR72.EXE activated at this point should be terminated prior to starting the DBMS installation procedure.

```
$ DBM$TMP = "SYS$SCRATCH:'F$UNIQUE()' .TMP
$ DEFINE /NOLOG /USER DBM$TMP 'DBM$TMP
$ ANALYZE /SYSTEM
    SET OUTPUT DBM$TMP
    SHOW PROCESS /CHANNELS ALL
    EXIT
$ SEARCH /OUTPUT='DBM$TMP' 'DBM$TMP';-1 DBMSHR72.EXE,"PID:"
$ SEARCH 'DBM$TMP' DBMSHR72.EXE /WINDOW=(1,0)
$ DELETE /NOLOG 'DBM$TMP';*{text}
```

In the following example, the process 2729F16D named "FOO\$SERVER" has the image DBMSHR72.EXE activated even after RMONSTOP72.COM has been executed and this process should be terminated prior to starting the DBM installation procedure:

```
$ @SYS$STARTUP:RMONSTOP72.COM
.
.
.
$ @FIND_DBMSHR72_PROC.COM

OpenVMS system analyzer

Process index: 016D   Name: FOO$SERVER   Extended PID: 2729F16D
0240 7FEF4460 8384F300 $1$DGA2:[VMS$COMMON.SYSLIB]DBMSHR72.EXE;722
$
```

4.2 Syntax for Prestart Transaction Timeout

The ability to define a timeout value for a prestarted transaction has been available on Oracle CODASYL DBMS since version V7.1, however the usage was never documented.

To define a timeout value for a prestarted transaction, enter the syntax:

```
$ DBO/MODIFY/TRANSACTION=PRESTART=(TIMEOUT=nnn) <database name>
```

where, 'nnn' is a value (range 0:3600) in seconds. This value represents the number of seconds to wait before aborting the prestarted transaction. Timing out the prestarted transaction may prevent snapshot file growth in environments where servers stay attached to the database with long periods of inactivity.

Additional functionality added to an update of V7.1 included an enhancement to force a process to obtain a new transaction sequence number (TSN) if the same TSN has been reused throughout the duration of the prestarted transaction timeout interval. This provides the ability for processes that constantly reuse TSNs to periodically obtain a new TSN, thus preventing excessive snapshot growth.

4.3 Patch Required When Using VMS V8.3 or V8.2–1 on Itanium

During qualification testing, a possible problem related to OpenVMS image activation on some Itanium systems was discovered. This problem could in some cases result in unexpected messages similar to the following:

```
%LIB-E-ACTIMAGE, error activating image  
-SYSTEM-F-NOPRIV, insufficient privilege or object protection violation
```

To avoid this problem, Oracle recommends that customers install one of the following architecture-specific patch kits (or subsequent replacement if superseded) for Itanium systems: VMS821I_SYS-V0400 (I64) VMS83I_SYS-V0200 (I64)

4.4 /LOCK_TIMEOUT Documentation Clarification

The /LOCK_TIMEOUT qualifier can be used on a number of DBO commands to specify how long an operation should wait for locks to be granted before timing out. These locks could be either the QUIET POINT lock or any other lock required for online operations.

Some Oracle CODASYL DBMS documentation suggests that there is a default value for this qualifier. However, this is not true.

If you specify the /LOCK_TIMEOUT qualifier, you must specify the number of seconds to wait:

```
$ DBO/BACKUP/AFTER/LOCK_TIMEOUT=5 . . .
```

will cause DBO to wait 5 seconds before returning a time out message.

4.5 Unexpected RCS termination

It has been observed in internal testing of Oracle Codasyl DBMS V7.2.2.0 that if the Record Cache Server (the RCS), terminates in an uncontrolled fashion this may under some conditions cause corruption of the database and/or the After Image Journal file.

When the RCS terminates the database is shut down and a message like the following is written to the monitor log:

```
6-DEC-2007 15:04:17.02 - Received Record Cache Server image termination from 22ED5144:1
- database name "device:[directory]database.R00;1" [device] (1200,487,0)
- abnormal Record Cache Server termination detected
- starting delete-process shutdown of database:
  - %DBM-F-RCSABORTED, record cache server process terminated abnormally
- sending process deletion to process 22ED10F9
- sending process deletion to process 22ECED59
- sending process deletion to process 22EC0158
- sending process deletion to process 22EB9543 (AIJ Log server)
- database shutdown waiting for active users to terminate
```

A future attempt to roll forward the AIJ following a restore of a database backup might fail with a bugcheck dump if this problem has happened.

The only currently known situation where this problem has been observed is if the logical name `DBM$BIND_RCS_VALIDATE_SECS` is defined to some value and the logical name `DBM$BIND_RCS_LOG_FILE` at the same time is undefined or defined incorrectly.

To prevent this problem, Oracle recommends any customer using the Row Cache feature either avoid defining the logical name `DBM$BIND_RCS_VALIDATE_SECS`, or if this logical name needs to be defined, makes sure that `DBM$BIND_RCS_LOG_FILE` is correctly defined (i.e. defined with the `/SYSTEM` and `/EXECUTIVE` qualifiers and pointing to a valid file name in an existing directory on a cluster accessible device with sufficient free space). This recommendation applies to all versions of Oracle Codasyl DBMS.

4.6 Slight Relaxation Of VMS\$MEM_RESIDENT_USER Requirement

Bug 5859487

Previously, the VMS\$MEM_RESIDENT_USER identifier was required to open a database that had any row cache configured for resident memory even if no caches were enabled for the database.

This restriction has been relaxed. If the database is not enabled for row caches, the VMS\$MEM_RESIDENT_USER identifier is not required even if caches are defined for resident memory.

4.7 Patch Required When Using VMS V8.3 and Dedicated CPU Lock Manager

During qualification testing of Oracle CODASYL DBMS Release 7.2.1 on OpenVMS V8` systems, a problem with the use of Extended Lock Value Blocks and the OpenVMS Dedicated CPU Lock Manager feature was discovered.

To avoid this problem, Oracle strongly recommends that customers wishing to use DBMS and the OpenVMS Dedicated CPU Lock Manager feature with OpenVMS V8.3 install one of the following architecture-specific patch kit (or subsequent replacement if superseded) prior to using Oracle CODASYL DBMS Release 7.2.1, or later, on OpenVMS V8.3 systems:

- VMS83I_SYS-V0200 (I64)
- VMS83A_SYS-V0100 (Alpha)

4.8 VMS\$MEM_RESIDENT_USER Rights Identifier Required

Oracle CODASYL DBMS release 7.1 introduced additional privilege enforcement for the database or row cache qualifiers MEMORY_MAPPING=SYSTEM and LARGE_MEMORY. If a database utilizes any of these features then the user account that opens the database must be granted the VMS\$MEM_RESIDENT_USER rights identifier. Also, any process attempting to change these attributes, to convert, or restore a database with these attributes enabled must also hold the right.

Oracle recommends that the DBO/OPEN command be used when utilizing these features.

4.9 Features Not Yet Available for OpenVMS I64

The following features or capabilities or components are not currently available to run or are known to not run reliably on OpenVMS I64 with this Oracle CODASYL DBMS release.

- Oracle CODASYL DBMS ADA precompiler
- PL/I compiler and Oracle CODASYL DBMS PL/I precompiler

4.10 Expect Additional Memory Consumption

Due to the increased sizes of image files (especially on Integrity servers) and more aggressive buffering and caching schemes and larger I/O size defaults, you should expect to allocate additional page file quota, working set sizes and buffered I/O byte limit quota when using Oracle CODASYL DBMS release 7.2. In particular, when running on Integrity servers, a page file quota of perhaps three times larger may be required for some applications.

4.11 ILINK-E-INVOVRINI Error on I64

When linking an application with multiple modules, the following error message may be returned:

```
%ILINK-E-INVOVRINI, incompatible multiple initializations for overlaid section
  section: DBM$UWA_B
  module: M1
  file: DKA0:[BLD]M1.OBJ;1
  module: M2
  file: DKA0:[BLD]SYS.OLB;1
```

On I64 systems, you cannot have a program section that attempts to be initialized a subsequent time where the non-zero portions of the initializations do not match. This is a difference from OpenVMS Alpha and VAX systems where the linker permitted such initializations.

This can be seen when linking multiple FORTRAN DML modules, where some modules use the default (non-stream) UWA, and another uses a "naked" invoke, which only contributes an abbreviated contribution to the DBM\$UWA_B psect.

For example, A.FOR contains:

```
PROGRAM AFOR
  INVOKE (SUBSCHEMA = FORTRAN_SUBSCHEMA,
1      SCHEMA = PARTS,
2      DATABASE = PARTS)
  CALL BSUB()
  END
```

B.FOR contains:

```
SUBROUTINE BSUB
  INVOKE
  RETURN
  END
```

On VAX or ALPHA, the above code will link and run correctly, However, on I64, the linker will generate the following:

```
%ILINK-E-INVOVRINI, incompatible multiple initializations for overlaid section
  section: DBM$UWA_B
  module: AFOR
  file: A.OBJ
  module: BSUB
  file: B.OBJ
```

4.12 Default Floating Point Format

The Itanium architecture has a 64-bit model and basic system functions similar to the Alpha chip. However, there are some implementation differences between the two platforms that might affect user-written applications.

One of the differences is the availability of hardware-supported floating-point formats. The Itanium architecture implements floating-point arithmetic in hardware using the IEEE floating-point formats, including IEEE single and IEEE double. The Alpha architecture supports both IEEE and VAX floating-point formats in hardware, and OpenVMS compilers generate code using the VAX formats by default, with options (on Alpha) to use IEEE formats. Irrespective of whether it was originally written for VAX or Alpha, an OpenVMS application that uses the default VAX floating-point formats needs to produce equivalent behavior on the Itanium architecture using IEEE formats at the lowest level.

- On OpenVMS VAX and OpenVMS Alpha, VAX float is the default. VAX format data is assumed and VAX floating instructions are used.
- On OpenVMS Alpha, you can specify the compiler option `/FLOAT=IEEE`. In this case, IEEE format data is assumed and IEEE floating instructions are used.
- On OpenVMS I64, IEEE float is the default. IEEE format data is assumed and IEEE floating instructions are used.
- On OpenVMS I64, you can specify the compiler option `/FLOAT=D_FLOAT` or `/FLOAT=G_FLOAT`.

When you compile an OpenVMS application that specifies an option to use VAX floating-point on the Itanium architecture, the compiler automatically generates code for converting floating-point formats. Whenever the application performs a sequence of arithmetic operations, this code does the following:

1. Converts VAX floating-point formats to either IEEE single or IEEE double floating-point formats.
2. Performs arithmetic operations in IEEE floating-point arithmetic.
3. Converts the resulting data from IEEE formats back to VAX formats.

Note that where no arithmetic operations are performed (VAX float fetches followed by stores), conversions will not occur. The code handles such situations as moves. VAX floating-point formats have the same number of bits and precision as their equivalent IEEE floating-point formats. For most applications, the conversion process will be transparent. In a few cases, arithmetic calculations might have different results because of the following differences between VAX and IEEE formats:

- Values of numbers represented
- Rounding rules
- Exception behavior

For more information, Oracle recommends reviewing the white paper "OpenVMS floating-point arithmetic on the Intel Itanium architecture" available from HP.

4.13 SYSTEM–F–INSFMEM Fatal Error With SHARED SYSTEM MEMORY or LARGE MEMORY Enabled in Galaxy Environment

When GALAXY support is enabled in an OpenVMS Galaxy environment, a *%SYSTEM–F–INSFMEM, insufficient dynamic memory error* message may be returned when mapping row caches or opening the database. One source of this problem specific to a Galaxy configuration is running out of Galaxy Shared Memory regions. For Galaxy systems, GLX_SHM_REG is the number of shared memory region structures configured into the Galaxy Management Database (GMDB).

While the default value of 64 regions (for OpenVMS versions through at least V7.3–1) might be adequate for some installations, sites using a larger number of databases or row caches when the SHARED MEMORY IS SYSTEM or LARGE MEMORY IS ENABLED features are enabled may find the default insufficient.

If a *%SYSTEM–F–INSFMEM, insufficient dynamic memory* error is returned when mapping record caches or opening databases, Oracle Corporation recommends that you increase the GLX_SHM_REG parameter by two times the sum of the number of row caches and number of databases that might be accessed in the Galaxy at one time. As the Galaxy shared memory region structures are not very large, setting this parameter to a higher than required value does not consume a significant amount of physical memory. It also may avoid a later reboot of the Galaxy environment. This parameter must be set on all nodes in the Galaxy.

Galaxy Reboot Required

Changing the GLX_SHM_REG system parameter requires that the OpenVMS Galaxy environment be booted from scratch. That is, all nodes in the Galaxy must be shut down and then the Galaxy reformed by starting each instance.

To enable Galaxy support, issue the command:

```
$ DBO/SET GALAXY/ENABLED <db>
```

To enable SYSTEM SHARED MEMORY, issue to command;

```
$ DBO/MODIFY/MEMORY_MAPPING=SYSTEM <db>
```

To enable LARGE MEMORY for record cache, issue the command:

```
$ DBO/CACHE/MODIFY/LARGE_MEMORY <db> <cache>
```

4.14 Oracle CODASYL DBMS and OpenVMS ODS-5 Volumes

The OpenVMS Version 7.2 release introduced an Extended File Specifications feature, which consists of two major components:

- A new, optional, volume structure, ODS-5, which provides support for file names that are longer and have a greater range of legal characters than in previous versions of OpenVMS.
- Support for "deep" directory trees.

ODS-5 was introduced primarily to provide enhanced file sharing capabilities for users of Advanced Server for OpenVMS 7.2 (formerly known as PATHWORKS for OpenVMS), as well as DCOM and JAVA applications.

In some cases, Oracle CODASYL DBMS performs its own file and directory name parsing and explicitly requires ODS-2 (the traditional OpenVMS volume structure) file and directory name conventions to be followed. Because of this knowledge, Oracle does not support any Oracle CODASYL DBMS database file components (including root files, storage area files, after-image journal files, record cache backing store files, database backup files, after-image journal backup files, and so forth) that utilize any non-ODS-2 file naming features. For this reason, Oracle recommends that Oracle CODASYL DBMS database components not be located on ODS-5 volumes.

Oracle CODASYL DBMS does support database file components on ODS-5 volumes provided that all of these files and directories strictly follow the ODS-2 file and directory name conventions. In particular, all file names must be specified entirely in uppercase and special characters in file or directory names are forbidden.

4.15 Carryover Locks and NOWAIT Transaction Clarification

In NOWAIT transactions, the BLAST (Blocking AST) mechanism cannot be used. For the blocking user to receive the BLAST signal, the requesting user must request the locked resource with WAIT (which a NOWAIT transaction does not do).

Oracle CODASYL DBMS defines a resource called NOWAIT, which is used to indicate that a NOWAIT transaction has been started. When a NOWAIT transaction starts, the user requests the NOWAIT resource. All other database users hold a lock on the NOWAIT resource so that when the NOWAIT transaction starts, all other users are notified with a NOWAIT BLAST.

The BLAST causes blocking users to release any carryover locks. There can be a delay before the transactions with carryover locks detect the presence of the NOWAIT transaction and release their carryover locks. You can detect this condition by examining the stall messages. If the "Waiting for NOWAIT signal (CW)" stall message appears frequently, the application is probably experiencing a decrease in performance, and you should consider disabling the carryover lock behavior.

4.16 Both Application and Oracle CODASYL DBMS Using SYS\$HIBER

In application processes that use Oracle CODASYL DBMS and the \$HIBER system service (possibly by RTL routines such as LIB\$WAIT), it is important that the application ensures that the event being waited for has actually occurred. Oracle CODASYL DBMS uses \$HIBER/\$WAKE sequences for interprocess communications particularly when the ALS (AIJ log server) feature is enabled.

The Oracle CODASYL DBMS use of the \$WAKE system service can interfere with other users of \$HIBER (such as the routine LIB\$WAIT) that do not check for event completion, possibly causing a \$HIBER to be unexpectedly resumed without waiting at all.

To avoid these situations, consider altering the application to use a code sequence that avoids continuing without a check for the operation (such as a delay or a timer firing) being complete.

The following pseudo-code shows one example of how a flag can be used to indicate that a timed-wait has completed correctly. The wait does not complete until the timer has actually fired and set TIMER_FLAG to TRUE. This code relies on ASTs being enabled.

```
ROUTINE TIMER_WAIT:
BEGIN
! Clear the timer flag
TIMER_FLAG = FALSE

! Schedule an AST for sometime in the future
STAT = SYS$SETIMR (TIMADR = DELTATIME, ASTRTN = TIMER_AST)
IF STAT <> SS$NORMAL
THEN BEGIN
LIB$SIGNAL (STAT)
END

! Hibernate. When the $HIBER completes, check to make
! sure that TIMER_FLAG is set indicating that the wait
! has finished.
WHILE TIMER_FLAG = FALSE
DO BEGIN
SYS$HIBER()
END
END

ROUTINE TIMER_AST:
BEGIN
! Set the flag indicating that the timer has expired
TIMER_FLAG = TRUE

! Wake the main-line code
STAT = SYS$WAKE ( )
IF STAT <> SS$NORMAL
THEN BEGIN
LIB$SIGNAL (STAT)
END
END
```

In OpenVMS V7.2, the LIB\$WAIT routine has been enhanced through the FLAGS argument (with the LIB\$K_NOWAKE flag set) to allow an alternate wait scheme (using the \$SYNCH system service) that can

avoid potential problems with multiple code sequences using the \$HIBER system service.

4.17 Row Cache Not Allowed While Hot Standby Replication is Active

The row cache feature may not be enabled on a Hot Standby database while replication is active. The Hot Standby feature will not start if row cache is enabled.

A new command qualifier, `/CACHE=NOENABLED`, has been added to the `DBO/OPEN` command. To open the Hot Standby database prior to starting replication, use the `/CACHE=NOENABLED` qualifier on the `DBO/OPEN` command.

4.18 Exclusive Access Transactions May Deadlock with RCS Process

If a record is frequently accessed by long running transactions that request read/write access, reserving the record for exclusive update, and if the record has one or more indexes, you may experience deadlocks between the user process and the row cache server (RCS) process.

There are at least three suggested workarounds to this problem:

1. Reserve the record for CONCURRENT UPDATE.
 2. Close the database and disable row cache for the duration of the exclusive transaction
 3. Change the checkpoint interval for the RCS process to a time longer than the time required to complete the batch job and then trigger a checkpoint just before the batch job starts. Set the interval back to a smaller interval after the checkpoint completes.
-

Chapter 5

New Features and Corrections in Previous Releases

[| Contents](#)